Perimeter and Coherence According to McCammond and Wise

Peter Brinkmann

Abstract

In [MW97], McCammond and Wise introduce the perimeter of 2-complexes and use it to obtain a sufficient criterion for coherence. We present a new exposition of their main result, as well as some applications.

1 Introduction

A group is called *coherent* if every finitely generated subgroup is finitely presented. Coherence is a commensurability invariant, i. e., if $H < G$ is a subgroup of finite index, then G is coherent if and only if H is coherent.

Free groups and surface groups are elementary examples of coherent groups. Moreover, fundamental groups of 3-manifolds are coherent (see [Sco73]), as are mapping tori of free group automorphisms (see [FH97]).

The proofs of these statements heavily use specific properties of the groups in question. Recently, however, Wise and McCammond have developed a more general approach to coherence that applies to a wide range of examples, such as certain one-relator groups, some small-cancellation groups, etc. (see [MW97]). The first four sections of this paper contain a new exhibition of their sufficient criterion for coherence, and the remaining sections contain some applications of this criterion.

It is my pleasure to thank my advisor S. M. Gersten for many helpful comments and discussions.

2 Conventions and definitions

All complexes and maps between complexes are assumed to be combinatorial. $f^{(n)} \in X$ denotes an n-cell of X. Deviating from standard terminology, cells of complexes are closed, and characteristic maps extend to the boundary of a cell. In particular, for a 2-cell f of X, the domain of the characteristic map $\chi_f : D \to X$ is a polygon.

For a 2-complex X, let X_S denote its stellar subdivision. Observe that the 2-cells of X_S are triangles. For a map $\Phi : Y \to X$ between 2-complexes, let Φ_* denote the induced map at the π_1-level and let $\Phi_S : Y_S \to X_S$ denote the induced map between the stellar subdivisions of X and Y. For an edge e of X, the *star* $St(e)$ of e is the collection of triangles of X_S that are adjacent to the edge of X_S corresponding to e.

2.1 Definition

A *weight function* on a 2-complex X is a map $w : \{\text{triangles of } X_S\} \to \mathbf{R}$. The weight $w(f)$ of a 2-cell f is the sum of the weights of the triangles in its stellar subdivision. The *standard weight function* assigns the weight 1 to every triangle in X_S. Unless stated otherwise, weight functions in this paper will take values in the nonnegative integers.

Given a weight function on X and a map $\Phi : Y \to X$, the *missing weight* of an edge e of Y is the sum of the weights of those triangles in the star of $\Phi(e)$ that are not contained in the image of the star of e. The *missing weight* of Φ is the sum of the missing weights of the edges of Y. In formulas, we have the double sum

$$M(\Phi) = \sum_{e^{(1)} \in Y} \; \sum_{f^{(2)} \in St(\Phi_S(e)) \setminus \Phi_S(St(e))} w(f).$$

The missing weight of a subcomplex $Y \subset X$ is the missing weight of the inclusion map.

Remark: We use the term “missing weight” instead of “perimeter” since the latter has caused some confusion.

Example: If X is a combinatorial 2-manifold without boundary (equipped with the standard weight function), then the missing weight of each 1-cell of X equals 2. If X is the 2-skeleton of the usual cubulation of $\mathbf{R}^3$, then the missing weight of each 1-cell is 4, and the missing weight of each 2-cell is 12.

A 2-cell f of X with characteristic map $\chi_f : D \to X$ is said to have *exponent* n if there is a closed path w in $X^{(1)}$ such that $\chi_f \mid_{\partial D}$ spells the word w^n and if n is maximal. The polygon D is symmetric under rotations by $\frac{2\pi}{n}$.

The *packet* $\tilde{f}$ of f is the complex obtained by taking n copies of D and gluing them along the boundary. There is a well defined map $\chi_{\tilde{f}} : \tilde{f} \to X$ whose restriction to the ith copy of D is χ_f precomposed with a rotation by $\frac{2\pi}{n}$.

A map $\Phi : Y \to X$ is said to be *packed* if for every 2-cell f of X and every lift γ of χ_f to Y there exists a lift of $\chi_{\tilde{f}}$ to Y extending γ. This condition is vacuous for any 2-cell whose characteristic map does not factor through Y. By attaching 2-cells to Y and extending Φ appropriately, we can assume that Φ is packed. Note that for nonnegative weights this does not increase the missing weight and that it does not change anything at the π_1-level.

2.2 Example

Let X be the presentation complex of $< x \mid x^n >$. X has exactly one 2-cell with characteristic map $\chi_f : D \to X$. Let $\Phi_1 : X^{(1)} \to X$ denote the inclusion of the 1-skeleton, and let Φ_2 be the restriction of χ_f to the boundary of D.

Then the missing weight of Φ_1 is exactly the weight of f, and packing Φ_1 yields a map of missing weight 0. The missing weight of Φ_2 is $n \cdot w(f)$, extending Φ_2 to the interior of D reduces the missing weight by $w(f)$, and packing Φ_2 yields a map of missing weight 0 since each copy of D reduces the missing weight by $w(f)$.

3 Reductions

Roughly speaking, the missing weight of a map $\Phi : Y \to X$ will play the role of a complexity function measuring to what extent Φ fails to be π_1-injective.

In order to show that a finitely generated subgroup $H < \pi_1 X$ is finitely presented, we will start with some finite complex Y and a map $\Phi : Y \to X$ such that $\Phi_*(\pi_1 Y) = H$. If Φ_* has nontrivial kernel, we will attach a suitable 2-cell to Y in order to make the kernel smaller. If we can do this in such a way that the missing weight goes down with each reduction of the kernel,

this process terminates after finitely many steps, proving that H is finitely presented. If this works for any finitely generated $H < \pi_1 X$, it follows that $\pi_1 X$ is coherent, and we will give sufficient conditions for this to happen.

3.1 Definition

Fix a packed map $\Phi : Y \to X$ and a 2-cell f of X with characteristic map $\chi_f : D \to X$. Let P be an interval subdivided into no more than $|\partial D|$ edges, where $|\partial D|$ denotes the number of faces of D.

A pair $(\rho : P \to D, \rho' : P \to Y)$ of immersed paths is called a *reduction of* Φ if the following conditions are satisfied:

1. ρ and ρ' fit into the following commutative diagram:

$$\begin{array}{ccc} P & \xrightarrow{\rho'} & Y \\ \rho \downarrow & \not\nearrow & \downarrow \Phi \\ D & \xrightarrow{\chi_f} & X \end{array}$$

2. There is no map $D \to Y$ that will fit into the diagram above.

If $|P| < |\partial D|$, the reduction (ρ, ρ') is said to be *incomplete*, in which case we define another immersed path $\sigma : S \to \partial D$ such that S is an interval subdivided into $|S| = |\partial D| - |P|$ edges and $\rho(P) \cup \sigma(S) = \partial D$. σ is called a *complement* of ρ, and it is unique up to orientation.

If $|P| = |\partial D|$, the reduction (ρ, ρ') is called *complete*. Finally, (ρ, ρ') is said to be *maximal* if there is no reduction (τ, τ') with the property that ρ and ρ' are proper subpaths of τ and τ', respectively.

Roughly speaking, reductions allow us to attach packets of 2-cells to Y, thus (under suitable hypotheses) reducing the missing weight of Φ. More precisely, we have the following

3.2 Construction

Fix a map $\Phi : Y \to X$ and a 2-cell f of X with characteristic map $\chi_f : D \to X$. If the pair $(\rho : P \to D, \rho' : P \to Y)$ is a reduction, then we construct a new map $\Phi^+ : Y^+ \to X$ in the following way:

If (ρ, ρ') is a complete reduction, then the endpoints of ρ are necessarily the same, even though the endpoints of ρ' in Y may not be equal. In this case, identify the endpoints of ρ' in Y, obtaining a new complex Y'. Φ factors through a map $\Phi' : Y' \to X$. Abusing notation, we still refer to the path $P \to Y \to Y'$ as ρ'. In the other case, i. e., if the endpoints of ρ' are already equal or if the reduction is incomplete, let $Y' = Y$ and $\Phi' = \Phi$. In any case, we have $M(\Phi) = M(\Phi')$.

Next, we define a complex Y'' by amalgamating D and Y' along P, i. e., $Y'' = Y' \cup_P D = (Y' \coprod D)/\sim$, where two elements $y \in Y'$ and $x \in D$ are equivalent precisely if there exists some $t \in P$ such that $\rho(t) = x$ and $\rho'(t) = y$. We can think of Y' and D as being subcomplexes of Y'', and we define a map $\Phi'' : Y'' \to X$ by $\Phi''|_{Y'} = \Phi'$ and $\Phi''|_D = \chi_f$.

Finally, we obtain $\Phi^+ : Y^+ \to X$ by packing $\Phi'' : Y'' \to X$. Note that $\Phi^+_*(\pi_1 Y^+) = \Phi_*(\pi_1 Y)$.

3.3 Lemma

Let $\Phi : Y \to X$ be a packed map, and let f be a 2-cell of X with characteristic map $\chi_f : D \to X$. If $(\rho : P \to D, \rho' : P \to Y)$ is a maximal incomplete reduction, then

$$M(\Phi^+) = M(\Phi) + M(\chi_f \circ \sigma) - n \cdot w(f),$$

where n is the exponent of f, σ is a complement of ρ and Φ^+ is the map constructed in 3.2.

Proof: Choose a word w in $X^{(1)}$ such that $\chi_f|_{\partial D}$ spells the word w^n, where n is the exponent of f. If $\rho' : P \to Y$ has a closed loop as a subpath, say ρ'', then $\Phi \circ \rho''$ necessarily spells a power of a conjugate of w in $X^{(1)}$ since $\Phi : Y \to X$ is combinatorial. This implies that the reduction ρ can be extended, which contradicts the maximality assumption. Hence, ρ' contains no closed loop.

Since ρ' contains no closed loop, the packet $\tilde{f}$ is a subcomplex of Y^+, and we have $M(\Phi^+|_{Y^+ \setminus \{\text{interior of 2-cells of } \tilde{f}\}}) = M(\Phi) + M(\chi_f \circ \sigma)$. Since there is no map $D \to Y$ that will fit into the diagram in definition 3.1, gluing back a 2-cell reduces the missing weight exactly by $w(f)$ (cf. example 2.2). $\square$

3.4 Lemma

Let $\Phi : Y \to X$ be a packed map, and let f be a 2-cell of X with characteristic map $\chi_f : D \to X$. If $(\rho : P \to D, \rho' : P \to Y)$ is a complete reduction, then

$$M(\Phi^+) \leq M(\Phi) - w(f),$$

where Φ^+ is the map constructed in 3.2.

Proof: Since ρ is complete, all the 1-cells in Y^+ correspond to 1-cells in Y, which implies that $M(\Phi) = M(\Phi^+|_{Y^+ \setminus \{\text{interior of 2-cells of } \tilde{f}\}})$. Since there is no map $D \to Y$ that will fit into the commutative diagram in definition 3.1, gluing the first 2-cell of $\tilde{f}$ back reduces the missing weight by $w(f)$, which proves the claim. Note that (as opposed to the previous proof) ρ' may contain closed loops, in which case gluing back subsequent 2-cells may not result in a further reduction of missing weight. $\square$

4 A sufficient condition for coherence

Let $\alpha : Q \to X$ be a contractible immersed loop in $X^{(1)}$. Then there exists a sequence of paths $\alpha_0 = \alpha, \cdots, \alpha_k = const$ such that α_{i+1} is obtained from α_i either by tightening, i. e., removing all subpaths of the form $e\bar{e}$, or by a homotopy across a 2-cell in the following way: A homotopy across a 2-cell f with characteristic map $\chi_f : D \to X$ uses a maximal reduction $(\rho : P \to D, \rho' :\to Q)$ such that

$$\begin{array}{ccc} P & \xrightarrow{\rho'} & Q \\ \rho \downarrow & & \downarrow \alpha_i \\ D & \xrightarrow{\chi_f} & X \end{array}$$

commutes. Note that we can think of P as being a subcomplex of Q. If ρ is complete, then α_{i+1} is constructed from α_i by removing P from Q. If ρ is incomplete, then we replace the subpath $\chi_f \circ \rho$ with its complement $\chi_f \circ \sigma$ (notation as in definition 3.1; we have to choose the correct the orientation of σ in order to obtain a continuous path).

4.1 Definition

The complex X is said to have the *path reduction property* if for each contractible immersed loop α there is a sequence $\alpha_0 = \alpha, \cdots, \alpha_k = const$ as above such that for every incomplete reduction (ρ, ρ') involving a 2-cell f we have the inequality

$$M(\chi_f \circ \sigma) \leq n \cdot w(f),$$

where n is the exponent of f and σ is a complement of ρ (cf. previous paragraph).

4.2 Theorem

Let X be a 2-complex enjoying the path reduction property for some weight function w. If the missing weight (with respect to w) of each edge of X is finite and if the weight of each 2-cell of X is strictly positive, then $\pi_1 X$ is coherent.

Proof: Let $H < \pi_1 X$ be a finitely generated subgroup. We want to show that H is finitely presented. There exists a finite 1-complex Y and a map $\Phi : Y \to X$ such that $\Phi_*(\pi_1 Y) = H$ and Φ is an immersion on the 1-skeleton (see [Sta83] for details).

If Φ is not π_1-injective, then there exists an essential loop β in Y such that $\alpha = \Phi \circ \beta$ is contractible in X. Choose a sequence $\alpha_0 = \alpha, \cdots, \alpha_k = const$ as above, using the path reduction property of X. The idea is to add finitely many 2-cells to Y in such a way that the loops $\alpha_0, \cdots, \alpha_k$ and the homotopies between them lift to the new complex.

Assume inductively that the loops $\alpha_0, \cdots, \alpha_i$ and the homotopies between them have already been lifted to $\beta_0 = \beta, \cdots, \beta_i$ in Y. If α_{i+1} is obtained from α_i by tightening, then α_{i+1} lifts to Y since Φ is an immersion on the 1-skeleton.

If a homotopy across a 2-cell fails to lift to Y, then the corresponding reduction $(\rho : P \to D, \rho' : P \to Q)$ of $\alpha_i : Q \to X$ gives rise to a reduction of Φ in the following way: The path α_i in the commutative diagram

$$\begin{array}{ccc} P & \xrightarrow{\rho'} & Q \\ \rho \downarrow & & \downarrow \alpha_i \\ D & \xrightarrow{\chi_f} & X \end{array}$$

factors through Y by our inductive hypothesis, inducing a reduction $(\rho, \beta_i \circ \rho')$ of Φ:

$$\begin{array}{ccccc} P & \xrightarrow{\rho'} & Q & \xrightarrow{\beta_i} & Y \\ \rho \downarrow & & & & \downarrow \Phi \\ D & & \xrightarrow{\chi_f} & & X \end{array}$$

Extend this to a maximal reduction of Φ and form the map $\Phi^+ : Y^+ \to X$ as in 3.2. Now the loop α_{i+1} lifts to Y^+, as does the homotopy taking α_i to α_{i+1}.

If the reduction is incomplete, then lemma 3.3 implies that $M(\Phi^+) = M(\Phi) + M(\chi_f \circ \sigma) - n \cdot w(f)$, and the path reduction property guarantees that $M(\Phi^+) \leq M(\Phi)$. If the reduction is complete, then, by lemma 3.4, $M(\Phi^+) \leq M(\Phi) - w(f) < M(\Phi)$ since $w(f) > 0$ by assumption. Φ^+ may not be an immersion on the 1-skeleton, but we can correct this by folding edges (see [Sta83]) without increasing the missing weight, which completes the inductive step.

Observe that we cannot arrive at a constant path unless at least one of the reductions of Φ is complete, so we end up with a finite complex Y' and a map $\Phi' : Y' \to X$ such that $\Phi'_*(\pi_1 Y') = H$, Φ' is an immersion on the 1-skeleton, and $M(\Phi') < M(\Phi)$.

Since we are using nonnegative integer weights, we can only repeat this process finitely many times before we arrive at a π_1-injective map $\Phi'' : Y'' \to X$ with finite domain Y'', which implies that H is finitely presented. Since H was arbitrary, this implies that $\pi_1 X$ is coherent. $\Box$

5 Applications

The applications listed in this section can be found in [MW97]; we give a unified approach to them using matchings in bipartite graphs. We will use the following concepts and results graph theory (see [Wes96], sec. 3.1 for details):

5.1 Matchings in bipartite graphs

All graphs are assumed to be finite. A graph G is called *bipartite* if its vertex set V can be expressed as the disjoint union of two nonempty sets V_1, V_2 such that every edge of G connects an element of V_1 to an element of V_2. In this case, a *matching of* V_1 *into* V_2 is a set M of pairwise disjoint edges of G such that every vertex in V_1 is contained in some edge in M.

A bipartite graph G with bipartition V_1, V_2 is said to have the *matching property with respect to* V_1 if for any subset $S \subset V_1$ the number of vertices in V_2 adjacent to S is at least the number of elements of S. The following theorem holds (see [Hal35, Wes96]):

5.2 Hall's matching condition

If G is a bipartite graph with bipartition V_1, V_2, then G has a matching of V_1 into V_2 if and only if G has the matching property with respect to V_1. $\square$

We now generalize Hall's theorem. Given a function m that assigns a positive integer to every vertex in V_1, we call a set M of edges an *m-matching* if each element of V_2 is contained in at most one element of M, and if for every vertex v in V_1, M contains exactly $m(v)$ edges emanating from v. We call $m(v)$ the *multiplicity* of v.

The graph G is said to have the *m-matching property* if for every subset W of V_1, the number of elements of V_2 adjacent to elements of W is at least as large as the sum of the multiplicities of the elements of W.

5.3 The m-matching theorem

Let G be a bipartite graph with bipartition V_1, V_2, and let m be a multiplicity function on V_1. Then G allows an m-matching if and only if G has the m-matching property.

Proof: Let G' be a graph whose vertex set is the disjoint union of V_2 and V_1', where V_1' contains vertices $v_1^x \cdots v_{m(x)}^x$ for every $x \in V_1$. For every edge (x, v) connecting $x \in V_1$ and $v \in V_2$ we choose edges $(v_1^x, v) \cdots (v_{m(x)}^x, v)$.

Now G' has the matching property if and only if G has the m-matching property, and G' has a matching of V_1' into V_2 if and only if G has an m-matching, so the claim follows from Hall's matching theorem. $\square$

6 Multiplicities, matchings and coherence

Let $\mathcal{P} =< X \mid R >$ be a finite presentation. The *incidence graph* of $\mathcal{P}$ is the bipartite graph G with vertex set $V = R \sqcup X$ such that $r \in R$ and $x \in X$ are connected by an edge if and only if x occurs in the spelling of r (relators are assumed to be cyclically reduced). Given some multiplicity function m on R, $\mathcal{P}$ is said to have the m-matching property if G has the m-matching property.

The following theorem is the main result of this section. It shows that, for the right choice of multiplicities, the m-matching property implies coherence. This theorem unifies the ideas behind several theorems in [MW97].

6.1 Theorem

Let the group G be given by the finite presentation $\mathcal{P}$ satisfying the m-matching property for some multiplicity function m. If we can reduce any word representing 1 *to the empty word by cyclic reduction or by replacing a subword of a relator r by its complement σ in such a way that*

$$n \cdot m(r) \geq |\sigma|, \tag{1}$$

then G is coherent. Here n is the exponent of r.

Proof: Let G be the incidence graph of $\mathcal{P} =< X \mid R >$, and let Y be the presentation complex of $\mathcal{P}$. By hypothesis, G has an m-matching M. Pick some edge in M, connecting some $r \in R$ to some $x \in X$. Now consider the 2-cell f of Y corresponding to r. At least one of the triangles in the stellar subdivision of f is adjacent to the 1-cell corresponding to x. Assign the weight one to one such triangle. Repeat this for all edges in M, and assign the weight zero to all remaining triangles.

Since each element of X belongs to at most one edge in M, the missing weight of each edge is at most one, so we have $|\sigma| \geq M(\sigma)$ for any path σ. Moreover, the weight of each 2-cell is exactly the multiplicity of the corresponding relator. Since inequality (1) holds, Y has the path reduction property. Since the weights of all 2-cells are positive, this shows that G is coherent. $\square$

When applying this theorem to a presentation $\mathcal{P}$, one typically begins with a multiplicity function m that satisfies (1), then one checks whether $\mathcal{P}$

has the m-matching property. Hence it is advantageous to choose m as small as possible. We list some reasonable choices for certain classes of presentations.

6.1.1 Dehn presentations

A presentation $\mathcal{P}$ is said to be a *Dehn presentation* if Dehn's algorithm solves the word problem for $\mathcal{P}$ (Dehn's algorithm solves the word problem for $\mathcal{P}$ if any cyclically reduced word representing the identity element contains a subword u of a relator r such that $|u| > \frac{|r|}{2}$).

For a Dehn presentation $\mathcal{P}$ and a relator r of $\mathcal{P}$, let n be the exponent of r and let w be a cyclically reduced word w such that $r = w^n$. We define $m(r) = \lfloor\frac{|w|-1}{2}\rfloor$ if $n = 1$, $m(r) = \lfloor\frac{|w|}{2}\rfloor$ if $n = 2$, and $m(r) = \lfloor\frac{|w|+1}{2}\rfloor$ if $n \geq 3$.

This choice of multiplicities satisfies (1) because it implies $n \cdot m(r) \geq \lfloor\frac{|r|-1}{2}\rfloor$. Hence, by theorem 6.1, a group given by a Dehn presentation is coherent if it has the matching property with respect to this choice of multiplicities.

6.1.2 Remark

Since presentations of type $C(4) - T(5)$ and $C(3) - T(7)$ define hyperbolic groups, it is natural to ask whether such presentations are Dehn presentations. The following two examples show that in general, the answer is no: Let $\mathcal{P}_1 =< a, b, c, d, r, s \mid ab\bar{a}r, \bar{r}\bar{b}cs, \bar{s}d\bar{c}\bar{d} >$ and $\mathcal{P}_2 =< a, b, c, d, r, s, t, u, v \mid abr, \bar{r}\bar{a}s, \bar{s}\bar{b}t, \bar{t}cu, \bar{u}dv, \bar{v}\bar{c}\bar{d} >$. $\mathcal{P}_1$ and $\mathcal{P}_2$ are of type $C(4)-T(5)$ and $C(3)-T(7)$, respectively, and both $\mathcal{P}_1$ and $\mathcal{P}_2$ are presentations of the fundamental group of the closed surface of genus two, so their star graph has exactly one cycle. This implies that – up to inversion and conjugacy – there is exactly one shortest boundary label w of van Kampen diagrams with exactly one interior vertex. It is easy to check that w cannot be shortened by means of a homotopy across a relator disc, which shows that $\mathcal{P}_1$ and $\mathcal{P}_2$ are not Dehn presentations.

6.1.3 Small cancellation conditions

A presentation is said to have property P if every piece has length one and no relator is a proper power (see [GS90]).

Since a nonempty cyclically reduced word representing 1 in a $C(6)$ presentation contains a complement of no more than three pieces, inequality (1) holds if we assign the multiplicity three to all the relators in a $C(6)-P$ presentation. Hence we recover a special case of theorem 9.4 of [MW97].

A nonempty cyclically reduced word representing 1 in a $C(4)-T(4)$ presentation contains a complement of no more than two pieces, so inequality (1) holds if we assign the multiplicity two to all the relators in a $C(4)-T(4)-P$ presentation. In this case we recover a special case of theorem 9.5 of [MW97].

6.1.4 λ-presentations

A presentation is said to be a λ-*presentation* if every nontrivial word representing the neutral element contains a subword of some relator r of length strictly greater than $(1-\lambda)|r|$. For example, Dehn presentations are $\frac{1}{2}$-presentations by definition, and B.B. Newman's spelling theorem shows that presentations of the form $< X \mid w^n >$ are $\frac{1}{n}$-presentations if $n \geq 2$.

For a relator r with exponent n in a λ-presentation, let k be the *largest* integer satisfying $k < \lambda|r|$. Then we choose $m(r)$ to be the *smallest* positive integer satisfying $n \cdot m(r) \geq k$. By construction, inequality (1) holds.

For example, if $\mathcal{P} =< X \mid w^n >$ for some $n \geq 2$, we have $k = |w|-1$, and $m(r) \geq \frac{|w|-1}{n}$. The incidence graph of $\mathcal{P}$ clearly has the m-matching property if the number of generators occuring in w is at least $\frac{|w|-1}{n}$.

In particular, the group given by $\mathcal{P}$ is coherent if $n \geq |w|-1$, so we recover theorem 7.3 in [MW97].

References

[FH97] Mark Feighn and Michael Handel. Mapping tori of free group automorphisms are coherent. preprint, 1997.

[GS90] S.M. Gersten and H.B. Short. Small cancellation theory and automatic groups. *Inventiones mathematicae*, 102:305–334, 1990.

[Hal35] P. Hall. On representation of subsets. *J. Lond. Math. Soc.*, 10:26–30, 1935.

[MW97] Jon McCammond and Daniel Wise. Coherent groups and the perimeter of 2-complexes. preprint, 1997.

[Sco73] P. Scott. Finitly generated 3-manifold groups are finitely presented. *J. London Math. Soc.*, 6(2):437–440, 1973.

[Sta83] J.R. Stallings. Topology of finite graphs. *Inv. Math.*, 71:551 – 565, 1983.

[Wes96] Douglas B. West. *Introduction to Graph Theory.* Prentice Hall, 1996.

Department of Mathematics, University of Utah
Salt Lake City, UT 84112, USA

E-mail: brinkman@math.utah.edu